RED BEAR: ALPHA ENGINE · DATA SHEET · HUMAN-IN-THE-LOOP REVIEW AND PREDICTIVE LEARNING

The system proposes. A named person decides.

Everything significant the system generates carries a confidence and a risk, reaches a person in proportion to both, and is kept with what it rested on and what rests on it.

KEY POINTS

- Everything significant the system generates carries a confidence and a risk, and is treated by band: held until a reviewer has acted, offered for review, or recorded and labelled unreviewed wherever it is shown
- The reviewer sees the evidence, the reasoning and what depends on the item, then approves, overrides with a rationale or escalates, and can ask why
- Banding sets how much human attention a proposal gets; no band lets a proposal approve, commit or change a record
- A decision ledger with lineage: overturn an entry and everything that depended on it surfaces for re-review; any past decision replays as it stood; the ledger exports as a signed package
- Learnings are proposed from outcomes, challenged by the analysts concerned, and used only once a person has accepted them; a learning is about a pattern and names no person

What it is

The product generates a great deal: criteria, fit assessments, elements of a company's ontology map, matched packs, baselines, suggested actions, learnings. This cluster is how people stay in charge of it. Each of those is a proposal with a confidence and a risk; the band it falls in fixes how much human attention it gets before anything relies on it; and it is recorded in a ledger with its lineage.

Learning follows the same rule. The system proposes what it has learned from outcomes, the analysts concerned challenge it and accept or reject it, and only an accepted learning is used.

The pain point

"The system produces more than anyone can check, and nobody can say which of its outputs a person actually looked at."

In our view a firm then has two bad options. Check everything, and its senior people become the queue. Check nothing in particular, and a generated figure travels into a committee paper with an authority nobody gave it. When one later proves wrong, nobody can say what rested on it, and the firm cannot show afterwards who looked at what, and when.

How it works

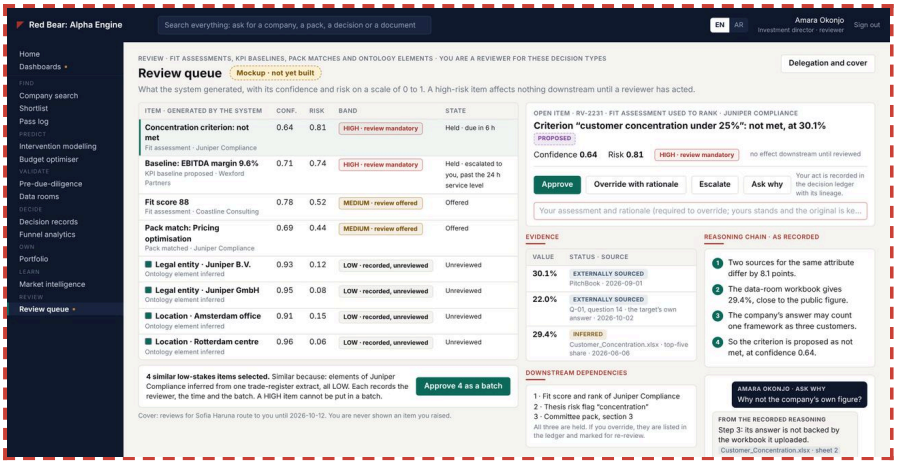
Confidence, risk and a band. A generated item is significant when a later step relies on it: a fit assessment used to rank, an inferred ontology element, a matched pack, a proposed baseline, an early warning's suggested action, a learning. Each carries a confidence score and a risk score. A high-risk item is held and has no effect downstream until a reviewer has acted; a medium-risk item is offered for review; a low-risk item is recorded and labelled unreviewed wherever it is shown (FR-REV-001).

The review queue. Each item shows its scores, its evidence and reasoning chain, and what depends on it. The reviewer approves, overrides with a mandatory rationale, the original kept beside the reviewer's value, or escalates, and may ask why the system proposed what it did; the answer is the recorded reasoning with its evidence. Nobody reviews an item they raised (FR-REV-002). A review administrator sets delegation and cover for absence; a mandatory review carries a service level and is escalated to the next named reviewer when it passes (FR-REV-004). Similar low-stakes items can be approved as a batch, the basis of similarity shown first and each item recording the reviewer; a high-risk item cannot be batched (FR-REV-005).

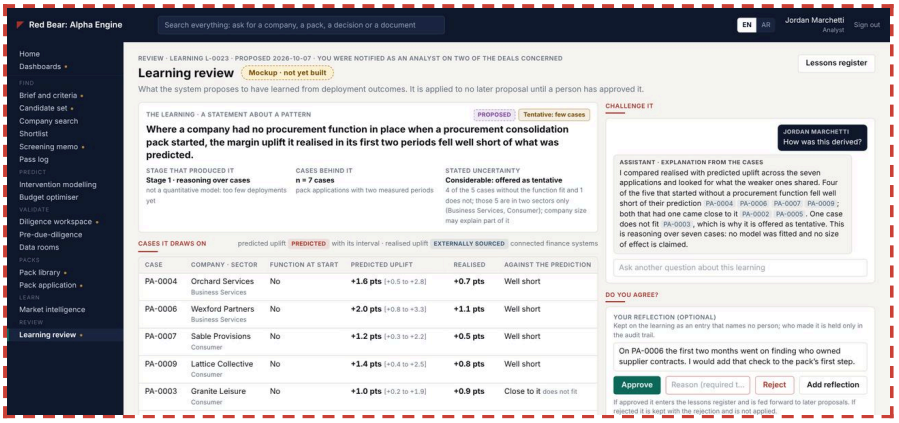
Governed, replayable, exportable. A governance administrator sets where the bands fall, the reviewer roles per decision type and the optional approvals. Settings are versioned, a change records who, when and why, and applies only to items raised afterwards; that high risk is reviewed is not configurable (FR-REV-006). Every generated item and every act of review sits in an append-only ledger with its inputs, evidence, scores, who acted and what depends on it; overturning an entry marks its dependants for re-review (FR-REV-003). Any past decision replays with its lineage and the settings then in force, altering nothing (FR-REV-007), and the ledger exports for a period or a deal as a signed package verifiable outside the system (FR-REV-011). The product reports where people most often override the system, by decision type, band and period, identifying no reviewer (FR-REV-008).

Learning a person confirms. A proposed learning is shown to the analysts concerned with the cases it draws on and its uncertainty. They may ask how it was derived, approve or reject it and add a reflection; it is applied to no later proposal until a person has approved it (FR-REV-009). Learning runs in two stages, reasoning over cases with stated uncertainty and then, once volume supports it, quantitative modelling as a versioned release; each learning states its stage and the number of cases behind it (FR-REV-010).

Screens



S-24 Review queue: The review queue: every item with its confidence, risk and band; the open item with its evidence, its recorded reasoning and what depends on it; and the reviewer's approve, override, escalate and ask why. Design-system mockup on synthetic data; replaced by the prototype render when the screen is built.



S-26 Learning review: Learning review: a proposed learning with the cases it draws on and its stated uncertainty, challenged in conversation and then approved or rejected by a person. Design-system mockup on synthetic data; replaced by the prototype render when the screen is built.

RETURN ON INVESTMENT

Attention where it matters. Senior time goes to the items whose risk warrants it, not to everything or to nothing. **An answer to "who looked at this?"** Every generated item shows whether a person reviewed it, and the proof replays. **A system that is corrected, visibly.** The measures this cluster carries are **the share of high-risk items reviewed before use**, which the ledger evidences rather than asserts; **time in queue against the service level**; and **where people most often override the system**, by decision type, band and period (FR-REV-008). All are recorded by the product from the first item raised; none is given a target here.

What it reads, and never reads

Reads. What the system generated, with its inputs, evidence, confidence and risk; the reviewers' acts and rationales; the governance settings in force; and the measured outcomes a learning draws on. **Never reads.** Message content. **Never computes.** A score, rank or count for a named reviewer or analyst; who acted is kept as a record, not as a measure. Minimum group sizes come from the platform's Privacy and aggregation service, and no lesson identifies who decided (RBP FR-PA-01, RBP FR-DL-04).

The boundaries

The system never decides. Banding sets how much human attention a proposal gets; it never lets a proposal approve, commit or change a record. Disagreement and throughput are reported as patterns by decision type, band and period, above a minimum group size, never per reviewer or per analyst. A learning is about a pattern and names no person; who approved it is held only in the audit trail.